

Revenue Maximization and Ex-Post Budget Constraints

CONSTANTINOS DASKALAKIS¹, M.I.T
 NIKHIL R. DEVANUR², Microsoft Research
 S. MATTHEW WEINBERG³, Princeton University

We consider the problem of a revenue-maximizing seller with m items for sale to n additive bidders with hard budget constraints, assuming that the seller has some prior distribution over bidder values and budgets. The prior may be correlated across items and budgets of the same bidder, but is assumed independent across bidders. We target mechanisms that are Bayesian Incentive Compatible, but that are *ex-post* Individually Rational and *ex-post* budget respecting. Virtually no such mechanisms are known that satisfy all these conditions and guarantee any revenue approximation, even with just a single item. We provide a computationally efficient mechanism that is a 3-approximation with respect to all BIC, *ex-post* IR, and *ex-post* budget respecting mechanisms. Note that the problem is NP-hard to approximate better than a factor of $16/15$, even in the case where the prior is a point mass [Chakrabarty and Goel 2010]. We further characterize the optimal mechanism in this setting, showing that it can be interpreted as a *distribution over virtual welfare maximizers*.

We prove our results by making use of a black-box reduction from mechanism to algorithm design developed by [Cai et al. 2013]. Our main technical contribution is a computationally efficient 3-approximation algorithm for the algorithmic problem that results by an application of their framework to this problem. The algorithmic problem has a mixed-sign objective and is NP-hard to optimize exactly, so it is surprising that a computationally efficient approximation is possible at all. In the case of a single item ($m = 1$), the algorithmic problem can be solved exactly via exhaustive search, leading to a computationally efficient exact algorithm and a stronger characterization of the optimal mechanism as a distribution over virtual value maximizers.

1. INTRODUCTION

Most of auction theory crucially depends on the assumption of quasi-linear utilities, that the utility is equal to valuation minus payments. This assumption fails when bidders are budget constrained.¹ Auctions with budget constrained bidders are commonplace, and prominent examples of this are ad-auctions and auctions for government licensing such as the FCC spectrum auction. An interesting example of budget constraint occurs in the auction for professional cricket players in the Indian Premier League: the league imposes a budget constraint on all the teams as a means of ensuring well balanced teams. Another source of budget constraints is what Che and Gale [1998] call the *moral hazard problem*: procurement is often delegated and budget constraints are imposed as a means of controlling the spend. A budget represents the bidder's *ability to pay*, in contrast to the valuation which represents his *willingness to pay*. For this reason, budgets may be more tangible and easier to estimate than

¹The terms financially constrained bidders or bidders with liquidity constraints are used synonymously.

¹ costis@csail.mit.edu, supported by a Microsoft Research Faculty Fellowship and NSF Awards CCF-0953960 (CAREER) and CCF-1101491.

² nikdev@microsoft.com.

³ sethwmw@cs.princeton.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© YYYY ACM. 0000-0000/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

valuations. It is therefore important to understand how budget constraints impact the design of auctions; this has been well established by now [Che and Gale 1998; Pai and Vohra 2014; Benoit and Krishna 2001; Laffont and Robert 1996; Maskin 2000; Malakhov and Vohra 2005; Che and Gale 2000; Bhattacharya et al. 2010].

The theory of auctions in the presence of budget constraints on bidders lags far behind the theory of auctions without budgets. For instance, consider the design of *optimal* (revenue maximizing) auctions that are Bayesian incentive compatible (BIC) and ex-post individually rational (IR). While Myerson [1981] gives a beautiful theory characterizing the optimal auction for any single parameter domain, no such characterization is known in the presence of private budgets (that could be correlated with the valuation). As a way to deal with this difficulty, previous papers have considered special cases and auctions with a subset of the desired properties. (See Section 1.2 for details.) We adopt the Computer Science approach of approximation, while incorporating all the desired properties. *The main result of this paper is a 3-approximation to the optimal auction in the class of auctions that are*

- *BIC,*
- *ex-post IR and*
- *ex-post budget respecting, with private budgets that could be correlated with the valuations,*

for multiple heterogeneous items and additive valuations. This is the first constant factor approximation for this class of auctions. Moreover, the computational problem, even without any incentive constraints is already NP-Hard to approximate within a ratio of 16/15 [Chakrabarty and Goel 2010]. This too suggests that an approximation is necessary.

1.1. Overview of Techniques

We prove our main result by making use of an algorithmic framework developed in [Cai et al. 2013]. The computational aspect of their framework provides a *black-box reduction* from a wide class of Bayesian mechanism design problems to problems of purely algorithm design. More specifically, they show that any α -approximation algorithm for a certain incentive-free algorithmic problem (induced by the mechanism design problem at hand) can be leveraged to find a BIC, IR mechanism that is also an α -approximation (to the optimal BIC, IR mechanism) in polynomial time. Significant further details on their reduction and how to employ it can be found in Section 2.1. After applying their framework to our problem, there is still the issue of solving the algorithmic problem that pops out of the reduction. This turns out to be essentially a (virtual) welfare maximization problem (without budgets), but where bidder types are somewhat involved. The optimization involves a mixed sign objective (i.e. the objective is a sum of several terms which can be positive or negative). Such optimization problems are typically solvable *exactly* in polynomial time or computationally hard to approximate within any finite factor, but rarely in between (due to the mixed signs in the objective). Interestingly, we obtain a 3-approximation for our mixed-sign objective problem despite the fact that it is NP-hard to optimize exactly. The design and analysis of our algorithm can be found in Section 3.

Cai et. al.'s framework also contains a structural result. We use it to show that the optimal auction in our setting is a *distribution* over virtual welfare maximizers. By this, we mean that the optimal mechanism maintains a distribution over n mappings, one mapping per bidder that maps types to virtual types, and, given a vector of reported types, it samples n mappings from this distribution, uses them to map the reported types to virtual types, and proceeds to choose an allocation that optimizes virtual welfare. Note that by virtual types in the previous sentence we do *not* mean the

specific virtual types as computed by Myerson’s virtual transformation, which aren’t even defined for multi-dimensional types, but just *some* virtual types that may or may not be the same as the true types. In particular, each mapping in the support of the mechanism’s distribution will take as input a type (which is an additive function with non-negative item values plus a non-negative budget), and output a virtual type without a budget constraint and whose valuation function is the sum of a budgeted-additive function² with non-negative item values (which depends on the input type in a very structured way) plus an additive function with possibly negative item values (which may be unstructured with respect to the input type). We provide a formal statement of this structural claim in Section 3 as well. Note that for the special case of a single item auction, this gives a particularly simple structure: the virtual types are now just a single (possibly negative) real number, which could be interpreted as a virtual value. The optimal auction simply maps reported types to virtual *values* and assigns the item to the bidder with the highest virtual value.

1.2. Related Work

The result that comes closest to characterizing the optimal auction is that of Pai and Vohra [2014]: they characterize the optimal budget respecting BIC auction for a single item. Their auction is implemented as an all-pay auction and is therefore not ex-post IR. They show that the optimal BIC, *interim* IR mechanism that respects budgets ex-post takes on a form similar to Myerson’s, but with additional *pooling* to enforce that no bidder is asked to pay more than her budget, while also maintaining that no bidder has incentive to underreport their budget. Earlier, Laffont and Robert [1996] and Maskin [2000] considered the case where valuations are private information but budgets are common knowledge and identical. Malakhov and Vohra [2005] study the setting where there are two bidders, one has a known budget constraint while the other does not. Che and Gale [2000] characterize the optimal pricing scheme for a single item with a single bidder, with private valuation and budget that could be correlated with each other. The limited special cases considered by these papers point to the difficulty of characterizing the optimal auction, which motivates the search for efficient approximations.

Another line of work ranks different auction formats by the revenue generated in the presence of budgets. Che and Gale [1998] compare first price, second price and all-pay auctions, while Benoit and Krishna [2001] compare sequential and simultaneous auctions.

In the computer science tradition, [Bhattacharya et al. 2010] give a 4-approximation for multiple items with additive valuations, but they assume that the budgets are publicly known, and the auction is not ex-post IR. Chawla et al. [2011] give a 2-approximation in a single parameter domain, but assume that the budgets are public. They also consider private budgets, where budgets and values are independently distributed, in single parameter matroid domains, and MHR Distributions, and give a $3(1 + e)$ -approximation. Finally, Cai et al. [2012], provide exactly optimal mechanisms for multiple items, additive valuations and private budgets, but their auctions are interim-IR. Once again, all these auctions make additional assumptions when compared to us.

Cai et al. [2013] give a general reduction from mechanism design to algorithm design, which we use for our results. For the special case of a single item auction with private budgets, we show that the algorithmic problem obtained through this reduc-

²A function $v(\cdot)$ is budgeted-additive if there exists a b such that $v(S) = \min\{b, \sum_{i \in S} v(\{i\})\}$ for all S . Note that this is different from an additive buyer with a budget, and that a budgeted-additive buyer indeed has quasi-linear utilities.

tion is quite easy to solve optimally, resulting in exactly optimal single-item auctions with budgets. However, when there are multiple items the resulting algorithmic problem becomes NP-Hard [Chakrabarty and Goel 2010]. We give a 3-approximation to this algorithmic problem which through the reduction gives a 3-approximately optimal multi-item auction with budgets. Recently, [Bhalgat et al. 2013] showed that (a weaker form of) the reduction of Cai et al. [2013] could be obtained using the simpler multiplicative weight update method instead of the ellipsoid algorithm used originally, and consider the variant of our setting where the items are *divisible*. The algorithmic problem in this case is once again easy. Daskalakis and Weinberg [2015] also use the reduction in [Cai et al. 2013] to design an auction for a non-linear objective, namely the *makespan* of an assignment of jobs to machines.

The auction design problem has also been considered in a *worst-case* model, as opposed to a Bayesian model. A standard framework is that of *competitive* auctions, where a bound is shown on the ratio of the revenue of an optimal auction to the revenue of the given auction on any instantiation of valuations and budgets. Borgs et al. [2005] and Abrams [2006] give constant competitive auctions for multi-unit auctions, under an assumption of *bidder dominance*, that the contribution of a single bidder to the total revenue is sufficiently small. [Devanur et al. 2013] give constant competitive auctions for single parameter downward-closed domains with a public, common budget constraint. Since the worst-case setting is decidedly more difficult than the Bayesian setting, these results are not comparable to ours. Another line of work considers the design of *Pareto-optimal* auctions: Dobzinski et al. [2008] characterize single item auctions that are Pareto-optimal, with public budgets and show an impossibility of a similar auction for private budgets. Goel et al. [2012] extend this auction to a more general poly-matroidal setting.

1.3. Conclusions and Future Work

The goals of revenue-optimality, ex-post individual rationality, and ex-post budget feasibility seem to be at odds with one another. This is highlighted by the fact that, prior to our work, no known auctions even approximately satisfied all three conditions, even with just a single item and private budgets that are independent of values. We provide a computationally efficient 3-approximation for the significantly more general case of auctions for multiple heterogeneous goods and additive bidders with private budgets that can be correlated with their values. While this model is already quite general compared to the previous state-of-the-art, it is an important direction to see if our results can be extended to more complex classes of bidder valuations, or to more complex constraints on feasible allocations. In particular, well studied classes of valuations such as gross substitutes would be interesting next steps.

2. PRELIMINARIES

We begin with formal definitions of the mechanism design problem we study. We then outline the reduction of Cai et al. [2013] (Section 2.1) and its implications (Section 2.2) for our problem. Finally we state a related problem (Section 2.3), the Generalized Assignment Problem, which we use in the design of our algorithm.

Bidders. There are n bidders, each with additive valuations over m items and a hard budget constraint. Specifically, bidder i has value v_{ij} for item j , value $\sum_{j \in S} v_{ij}$ for set S , and hard budget b_i . We denote by \vec{v}_i the vector of bidder i 's values for all m items. We denote by \mathcal{D}_i the joint distribution of (\vec{v}_i, b_i) . We denote by $\mathcal{D} = \times_i \mathcal{D}_i$ the joint distribution of all bidders' valuations and budgets.

Mechanisms. Our goal is to design *Bayesian Incentive Compatible* (BIC) mechanisms that are *ex-post Individually Rational* (IR) and that respect budgets ex-post. Formally, for a (randomized) mechanism M , we can denote by $x_{ij}^M(\vec{v}, \vec{b}, r)$ to be 1 if bidder i receives item j when the profile of values/budgets reported to M is (\vec{v}, \vec{b}) and the random seed used by M is r , or 0 otherwise. Similarly, we denote by $q_i^M(\vec{v}, \vec{b}, r)$ to be the price paid by bidder i in the same conditions. We can then define the interim allocation probability $\pi_{ij}^M(\vec{v}_i, b_i)$ to be the probability that bidder i receives item j when reporting (\vec{v}_i, b_i) over the randomness of other agent's (valuation, budget)s $(\vec{v}_{-i}, \vec{b}_{-i})$ being drawn from \mathcal{D}_{-i} , and any randomness in M . We can similarly define the interim price $p_i^M(\vec{v}_i, b_i)$ to be the expected payment made by bidder i over the same randomness. Formally, $\pi_{ij}^M(\vec{v}_i, b_i) = \mathbb{E}_{(\vec{v}_{-i}, \vec{b}_{-i}) \leftarrow \mathcal{D}_{-i}, r} [x_{ij}^M(\vec{v}_i; \vec{v}_{-i}, b_i; \vec{b}_{-i}, r)]$ and $p_i^M(\vec{v}_i, b_i) = \mathbb{E}_{(\vec{v}_{-i}, \vec{b}_{-i}) \leftarrow \mathcal{D}_{-i}, r} [q_i^M(\vec{v}_i; \vec{v}_{-i}, b_i; \vec{b}_{-i}, r)]$. Formal definitions of BIC, IR, and ex-post budgets are below.

Definition 2.1. (Bayesian Incentive Compatible) A mechanism M is BIC if for all bidders i , and types $(\vec{v}_i, b_i), (\vec{v}'_i, b'_i)$ the following holds:

$$\vec{v}_i \cdot \pi_i^M(\vec{v}_i, b_i) - p_i^M(\vec{v}_i, b_i) \geq \vec{v}_i \cdot \pi_i^M(\vec{v}'_i, b'_i) - p_i^M(\vec{v}'_i, b'_i).$$

A mechanism is said to be ϵ -BIC if for all bidders i , and types $(\vec{v}_i, b_i), (\vec{v}'_i, b'_i)$ the following holds:

$$\vec{v}_i \cdot \pi_i^M(\vec{v}_i, b_i) - p_i^M(\vec{v}_i, b_i) \geq \vec{v}_i \cdot \pi_i^M(\vec{v}'_i, b'_i) - p_i^M(\vec{v}'_i, b'_i) - \epsilon.$$

Definition 2.2. (Interim/Ex-Post Individually Rational) A mechanism M is interim IR if for all bidders i , and types (\vec{v}_i, b_i) the following holds:

$$\vec{v}_i \cdot \pi_i^M(\vec{v}_i, b_i) \geq p_i^M(\vec{v}_i, b_i).$$

Further, it is ex-post IR if for all bidders i , all profiles (\vec{v}, \vec{b}) and random seeds r , we have:

$$\vec{v}_i \cdot x_i^M(\vec{v}, \vec{b}, r) \geq q_i^M(\vec{v}, \vec{b}, r).$$

Definition 2.3. (Ex-Post Budget Respecting) A mechanism M respects budgets ex-post if for all type profiles (\vec{v}, \vec{b}) , all random seeds r , and all bidders i we have:

$$q_i^M(\vec{v}, \vec{b}, r) \leq b_i.$$

Definition 2.4. (No Positive Transfers) A mechanism M has no positive transfers if for all type profiles (\vec{v}, \vec{b}) , all random seeds r , and all bidders i we have:

$$q_i^M(\vec{v}, \vec{b}, r) \geq 0.$$

2.1. Reduction from Mechanism to Algorithm Design

In recent work, [Cai et al. 2013] provide an algorithmic framework for mechanism design, showing how to design mechanisms by solving purely algorithmic problems. We use this reduction to reduce our mechanism design problem to an algorithm design problem and show a 3-approximation to this algorithmic problem. In the rest of this section, we state the general formulations of the mechanism design and the corresponding algorithm design problems considered by Cai et al. [2013]. Then we give the precise statement of their reduction, and a structural characterization of the optimal mechanism obtained as a byproduct of their reduction. Finally we instantiate these to state the corresponding problems in our setting, and massage the resulting problems to simplify them.

[Cai et al. 2013] call the mechanism design problems of study $\text{BMeD}(\mathcal{F}, \mathcal{V}, \mathcal{O})$,³ where feasibility constraints \mathcal{F} , possible valuations \mathcal{V} , and optimization objective \mathcal{O} parameterize the problem. Formally, this problem is defined as:

BMeD($\mathcal{F}, \mathcal{V}, \mathcal{O}$):

INPUT: For each bidder $i \in [n]$, a finite set $T_i \subseteq \mathcal{V}$, and a distribution D_i over T_i , presented by explicitly listing all types in T_i and their corresponding probability.

OUTPUT: A feasible (selects an allocation in \mathcal{F} with probability 1), BIC, (interim) IR mechanism for bidders drawn from $D = \times_i D_i$.

GOAL: Find the mechanism that optimizes \mathcal{O} in expectation, with respect to all BIC, IR mechanisms (when bidders with types drawn from D play truthfully).

APPROXIMATION: An algorithm is said to be an (ϵ, α) -approximation if it finds an ϵ -BIC mechanism whose expected value of \mathcal{O} (when bidders drawn from D report truthfully) is at least $\alpha \cdot \text{OPT} - \epsilon$.

In our problem, the feasible allocations are those that award each item to at most one bidder. So we could denote the set of feasible allocations as $[m+1]^n$ (with the convention that selecting the allocation \vec{a} awards item j to bidder a_j if $a_j > 0$, or no one if $a_j = 0$). The possible bidder types are all additive functions over items (with non-negative multipliers), and non-negative budgets, which we could denote by \mathbb{R}_+^{m+1} . Our objective is revenue. To ensure that all feasible mechanisms are ex-post IR (note that their reduction only guarantees interim IR without extra work) and ex-post budget respecting, we will define the objective function REVENUE as follows. REVENUE takes as input a valuation profile (\vec{v}, \vec{b}) , an allocation \vec{x} (where $x_{ij} = 1$ iff bidder i is awarded item j), and a price vector \vec{p} . We define $\text{REVENUE}(\vec{v}, \vec{b}, \vec{x}, \vec{p}) = \sum_i p_i$, if $0 \leq p_i \leq \min\{b_i, \vec{v}_i \cdot \vec{x}_i\}$ for all i , or $\text{REVENUE}(\vec{v}, \vec{b}, \vec{x}, \vec{p}) = -\infty$ otherwise.

There is a subtle issue with respect to *why* we want to design mechanisms that respect budgets ex-post. Specifically, is it just because the designer wishes to offer this guarantee to the bidders, who have true quasi-linear preferences? If so, then this is exactly the setting we have described so far: the designer is constrained to select a mechanism that respects budgets ex-post, but bidders will still choose how to play as if they were quasi-linear. While this motivation is certainly mathematically interesting, it is also non-standard and perhaps unrealistic. Instead, the more common motivation is because bidders physically can't pay more than their budget, and would have utility $-\infty$ if asked to do so. In this case, the designer should actively *exploit* this to extract higher revenue. For example, if for all i , bidder i is awarded all the items and charged her budget with tiny probability ϵ/n , then the designer needn't worry about bidders overreporting their budget (as otherwise they'd get utility $-\infty$ with probability ϵ/n). Therefore, a mechanism can be made BIC in the latter case (while losing arbitrarily little revenue) iff for all $(\vec{v}_i, b_i), (\vec{v}'_i, b'_i)$ **with** $b'_i \leq b_i$, $\vec{v}_i \cdot \vec{\pi}_i^M(\vec{v}_i, b_i) - p_i^M(\vec{v}_i, b_i) \geq \vec{v}'_i \cdot \vec{\pi}_i^M(\vec{v}'_i, b'_i) - p_i^M(\vec{v}'_i, b'_i)$. Note that this is a relaxed condition of the former setting, which requires the inequality to hold **for all** b_i, b'_i . Fortunately, the Cai et. al. framework applies in both settings, and the resulting structure and algorithmic problem are exactly the same. So all of our theorems, exactly as stated, hold in both of the described settings.

Informally, the main result of [Cai et al. 2013] states that, for all $\mathcal{F}, \mathcal{V}, \mathcal{O}$, the problem $\text{BMeD}(\mathcal{F}, \mathcal{V}, \mathcal{O})$ can be solved in polynomial time with black-box access to a poly-time algorithm for a purely algorithmic problem that they call $\text{GOOP}(\mathcal{F}, \mathcal{V}, \mathcal{O})$.⁴

³BMeD stands for **B**ayesian **M**echanism **D**esign.

⁴GOOP stands for **G**eneralized **O**bjective **O**ptimization **P**roblem.

Below, \mathcal{V}^\times denotes the closure of \mathcal{V} under addition and (possibly negative) scalar multiplications (so for instance, $(\mathbb{R}_+^m)^\times = \mathbb{R}^m$).

GOOP($\mathcal{F}, \mathcal{V}, \mathcal{O}$):

INPUT: A type $t_i \in \mathcal{V}$, multiplier $m_i \in \mathbb{R}$, and virtual valuation function $g_i \in \mathcal{V}^\times$ for each $i \in [n]$.⁵

OUTPUT: An allocation $x \in \mathcal{F}$ and price vector $\vec{p} \in \mathbb{R}_+^n$.

GOAL: Find $\arg \max_{x \in \mathcal{F}, \vec{p}} \{\mathcal{O}(\vec{t}, x, \vec{p}) + \sum_i m_i p_i + \sum_i g_i(x)\}$.

APPROXIMATION: (x^*, \vec{p}^*) is said to be an α -approximation if $\mathcal{O}(\vec{t}, x^*, \vec{p}^*) + \sum_i m_i p_i^* + \sum_i g_i(x^*) \geq \alpha \cdot \arg \max_{x \in \mathcal{F}, \vec{p}} \{\mathcal{O}(\vec{t}, x, \vec{p}) + \sum_i m_i p_i + \sum_i g_i(x)\}$.

Further below we provide much more detail on the structure of the algorithmic focus of this paper, **GOOP**($[n+1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE}$), but we first conclude our discussion of the reduction we employ. The main result of [Cai et al. 2013] states that for all $\epsilon > 0$, an (ϵ, α) -approximation for **BMeD**($\mathcal{F}, \mathcal{V}, \mathcal{O}$) can be obtained from a poly-time α -approximation for **GOOP**($\mathcal{F}, \mathcal{V}, \mathcal{O}$). The additive error (and failure probability in the theorem statement) is due to a sampling procedure in the execution of the reduction. We provide a full statement of their main result below.⁶

THEOREM 2.5. (Theorem 4 of [Cai et al. 2013]) *For all $\mathcal{F}, \mathcal{V}, \mathcal{O}$, and $\epsilon > 0$, if there is a poly-time α -approximation algorithm, G , for **GOOP**($\mathcal{F}, \mathcal{V}, \mathcal{O}$), there is a poly-time (ϵ, α) -approximation algorithm for **BMeD**($\mathcal{F}, \mathcal{V}, \mathcal{O}$) as well. Specifically, if ℓ denotes the input length of a **BMeD**($\mathcal{F}, \mathcal{V}, \mathcal{O}$) instance, the algorithm runs in time $\text{poly}(\ell, 1/\epsilon)$, makes $\text{poly}(\ell, 1/\epsilon)$ black box calls to G on inputs of size $\text{poly}(\ell, 1/\epsilon)$, and succeeds with probability $1 - \exp(-\text{poly}(\ell, 1/\epsilon))$.*

[Cai et al. 2013] prove Theorem 2.5 above by considering a linear program that optimizes over the space of interim forms that are both truthful (that satisfy the linear constraints in Definitions 2.1 and 2.2), and feasible (those that correspond to an actual mechanism that selects an outcome $x \in \mathcal{F}$ on every profile with probability 1).⁷ Linear constraints enforcing that an interim form is BIC and interim IR can be written explicitly, but a computationally efficient separation oracle for the space of feasible interim forms is still required in order to solve the linear program. They show how to obtain such a separation oracle with black-box access to an algorithm that solves **GOOP**, and that this entire process preserves approximation as well.

[Cai et al. 2013] further provide a structural characterization of the space of all feasible mechanisms (truthful or not), leading to a structured implementation of whatever interim form is output by the LP. Specifically, they show that the extreme points of the space of feasible interim forms correspond to mechanisms that associate a virtual valuation function $g_i(t_i)(\cdot)$ and price multiplier $m_i(t_i)$ to each type $t_i \in T_i$, and then selects on profile (t_1, \dots, t_n) the allocation and price vector that solves **GOOP** on input $t_1, \dots, t_n, m_1(t_1), \dots, m_n(t_n), \sum_i g_i(t_i)(\cdot)$. They show further that solving the linear program explicitly finds a list of virtual valuation functions and multipliers whose resulting interim forms contain the optimal (truthful) interim form in their convex hull. Theorem 2.6 below captures the structural aspect of their result.

⁵For other applications, the inputs $g_i(\cdot)$ are sometimes called instead cost functions.

⁶The theorem statement is identical in content, but reworded for clarity and cleanliness.

⁷In fact, they need to work with a generalization of interim forms, called *implicit forms*, to accommodate non-additive valuations. But we describe their proof for additive valuations for clarity of exposition, and because it is relevant for our setting.

THEOREM 2.6. *(Implicit in [Cai et al. 2013]) For all BMeD instances, the optimal mechanism can be implemented as a distribution over generalized objective optimizers. Specifically, there exists a distribution Δ over mappings $(f_1^\delta, \dots, f_n^\delta)$. Each mapping f_i^δ takes types t_i in T_i to price multipliers $m_i^\delta(t_i) \in \mathbb{R}$ and virtual valuation functions $g_i^\delta(t_i)(\cdot) \in \mathcal{V}^\times$. The optimal mechanism first samples $(f_1^\delta, \dots, f_n^\delta)$ from Δ , and on profile \vec{t} , selects the outcome and price vector $\arg \max_{x \in \mathcal{F}, \vec{p}} \{O(\vec{t}, x, \vec{p}) + \sum_i m_i^\delta(t_i) \cdot p_i + \sum_i g_i^\delta(t_i)(x)\}$.*

In the section below, we provide further details surrounding instantiations of Theorems 2.5 and 2.6 as they pertain to the problem at hand.

2.2. Instantiations

The goal of this section is to provide more details of the instantiation of Theorems 2.5 and 2.6 to our setting, but not to provide proofs (for which we refer the reader to [Cai et al. 2013]). We begin by describing the linear program that the reduction of [Cai et al. 2013] would try to solve for our setting. Below, $F([n+1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE})$ denotes the space of interim forms of all feasible (not necessarily truthful) mechanisms. Specifically, $(O, \vec{\pi}, \vec{p}) \in F([n+1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE})$ if and only if there is a mechanism M that awards each item at most once on every profile, is ex-post IR and ex-post budget respecting, awards bidder i item j when she reports type t_i with probability exactly $\pi_{ij}(t_i)$ (w.r.t. all other bidders' types and the randomness in the mechanism) and charges bidder i price $p_i(t_i)$ in expectation (over all other bidders' types and the randomness in the mechanism), and whose expected revenue is exactly O . With this definition in mind, the linear program they solve is stated below.

Variables:

- O , denoting the expected revenue of the interim form found.
- $\pi_{ij}(t_i)$ for all bidders i , items j , types t_i , denoting the probability that bidder i receives item j when reporting type t_i .
- $p_i(t_i)$ for all bidders i and types t_i , denoting the expected price paid by bidder i when reporting type t_i .

Constraints:

- (1) $\sum_j \pi_{ij}(t_i) \cdot v_{ij}(t_i) - p_i(t_i) \geq \sum_j \pi_{ij}(t'_i) \cdot v_{ij}(t_i) - p_i(t'_i)$, for all bidders i and types t_i, t'_i , guaranteeing that the interim form corresponds to a BIC mechanism.
- (2) $\sum_j \pi_{ij}(t_i) \cdot v_{ij}(t_i) - p_i(t_i) \geq 0$, for all bidders i and types t_i , guaranteeing that the interim form corresponds to an interim IR mechanism.⁸
- (3) $(O, \vec{\pi}, \vec{p}) \in F([n+1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE})$, guaranteeing that the interim form corresponds to a feasible mechanism.

Maximizing:

- O , the expected revenue.

The solution to this LP is the interim form of the optimal mechanism. The LP can be solved in polynomial time, so long as we have a poly-time separation oracle for the space $F([n+1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE})$. [Cai et al. 2013] shows that this can be obtained

⁸Actually, this constraint is redundant as we will also enforce that the mechanism be ex-post IR to be considered feasible.

via an algorithm for the related GOOP problem, which we instantiate in our setting below.

Budgeted-Additive Virtual Welfare Maximization. As discussed above, in order to find (approximately) optimal mechanisms for our setting, we need to study the purely algorithmic problem $\text{GOOP}([n+1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE})$, which we pose formally below.

GOOP $([n+1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE})$:

INPUT: Values $v_{ij} \geq 0$ and virtual values $w_{ij} \in \mathbb{R}$ for all i, j . Budget $b_i \in \mathbb{R}_+$ and price multiplier $m_i \in \mathbb{R}$ for all i .

OUTPUT: An allocation $\vec{x} \in \{0, 1\}^{mn}$ and prices \vec{p} such that $\sum_i x_{ij} \leq 1$ for all j (each item awarded at most once), $\sum_j x_{ij} v_{ij} \geq p_i$ (ex-post IR), $p_i \leq b_i$ (ex-post budget respecting), and $p_i \geq 0$ (no positive transfers).

GOAL: Find $\arg \max_{\vec{x}, \vec{p}} \{ \sum_i (m_i + 1) p_i + \sum_{i,j} x_{ij} w_{ij} \}$ (virtual revenue plus virtual welfare).

Note that in the above formulation, we have folded cases where REVENUE evaluates to $-\infty$ into feasibility constraints on the output. We make two quick further observations about the structure of $\text{GOOP}([n+1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE})$, and call the reformulation **Budgeted-Additive Virtual Welfare Maximization (BAVWM)**. Also, for cleanliness, we will replace the input price multipliers m_i by $m_i - 1$ so that the term in the objective will be $\sum_i m_i p_i$. This is w.l.o.g. as each m_i could be any real number.

OBSERVATION 1. *If $m_i > 0$, the optimal choice for p_i is always $\min\{b_i, \sum_j x_{ij} v_{ij}\}$. If $m_i \leq 0$, the best choice for p_i is 0.*

OBSERVATION 2. *For all possible solutions (\vec{x}, \vec{p}) , the quality of (\vec{x}, \vec{p}) for the input instance $(\vec{v}, \vec{w}, \vec{b}, \vec{m})$ is the same as for the instance $(\vec{v}', \vec{w}, \vec{b}, \vec{m})$ where $v'_{ij} = \min\{v_{ij}, b_i\}$, for all i, j .*

In light of these, we may set all negative m_i to 0, and all v_{ij} to $\min\{v_{ij}, b_i\}$ without changing the problem, leading to the following reformulation.

Budgeted-Additive Virtual Welfare Maximization:

INPUT: Budget b_i for all agents. Values $v_{ij} \in [0, b_i]$ for all agents and items. Price multiplier $m_i \geq 0$ for all agents, and virtual value $w_{ij} \in \mathbb{R}$ for all agents and items.

OUTPUT: An allocation $\vec{x} \in \{0, 1\}^{mn}$ such that $\sum_i x_{ij} \leq 1$ for all j (each item awarded at most once).

GOAL: Find $\arg \max_{\vec{x}} \{ \sum_i (m_i \min\{b_i, \sum_j x_{ij} v_{ij}\} + \sum_{i,j} x_{ij} w_{ij}) \}$.

Note that in the above formulation, we no longer need to optimize over the price vector, due to Observation 1. The problem can now be interpreted as just a welfare maximization problem, where bidder i 's valuation function is the sum of a budgeted-additive function (with non-negative item values) and an additive function (with possibly negative item values). Also, note that we can re-formulate the above problem to remove the multipliers $(m_i)_i$ from the input and the objective, by incorporating them in the b_i 's and the v_{ij} 's. We choose to leave them in so that it is more transparent how the inputs to BAVWM are related to the types reported by the bidders of the mechanism output by the Cai et al. [2013] reduction.

2.3. The Generalized Assignment Problem

Our main technical result will make use of a rounding algorithm for the *Generalized Assignment Problem*. We give here a statement of the problem and a rounding theorem due to Shmoys and Tardos [Shmoys and Tardos 1993].

Generalized Assignment Problem:

INPUT: Processing times $p_{ij} \in \mathbb{R}_+$ and costs $c_{ij} \in \mathbb{R}$ for all machines i and jobs j , capacities T_i for all machines i .⁹

OUTPUT: An allocation $\vec{x} \in \{0, 1\}^{mn}$ of jobs to machines such that $\sum_i x_{ij} = 1$ for all j (each job is assigned) and $\sum_j x_{ij} p_{ij} \leq T_i$ (each machine processes at most its capacity).

GOAL: Find $\arg \max_{\vec{x}} \{\sum_{i,j} x_{ij} c_{ij}\}$ (total cost).¹⁰

Now, we provide an LP due to Shmoys and Tardos that outputs a fractional solution at least as good as OPT.

Variables:

— x_{ij} , for all machines i and jobs j , denoting the fraction of job j assigned to machine i .

Constraints:

- (1) $\sum_i x_{ij} = 1$, for all j , guaranteeing that every job is processed exactly once.
- (2) $\sum_j x_{ij} \leq T_i$, for all i , guaranteeing that no machine's capacity is violated.
- (3) $x_{ij} = 0$ if $p_{ij} > T_i$.

Maximizing:

— $\sum_{i,j} x_{ij} c_{ij}$, the total cost.

THEOREM 2.7. ([Shmoys and Tardos 1993]) *The optimal fractional solution to the above LP can be rounded in polynomial time to an integral solution such that:*

- (1) $\sum_i x_{ij} = 1$, for all j .
- (2) $\sum_j x_{ij} \leq 2T_i$, for all i .
- (3) $\sum_j x_{ij} c_{ij} \geq \text{OPT}$.

3. MAIN RESULTS

In Section 3.1 below, we provide our main computational result: a poly-time approximation algorithm for BAVWM, which implies a poly-time truthful mechanism for revenue maximization that respects ex-post IR and ex-post budget constraints. In Section 3.2, we detail the structure of the optimal mechanism in this setting, as well as our computationally efficient mechanism from Section 3.1.

3.1. Computational Results

In this section, we provide a poly-time 3-approximation for BAVWM. We begin by writing a LP relaxation, allowing the designer to award fractions of items as long as the total fraction awarded doesn't exceed 1. We split the fraction of item j awarded to bidder i into two parts, \bar{x}_{ij} and \hat{x}_{ij} . Let \bar{x}_{ij} denote the fraction of item j assigned to agent i before exceeding b_i . And let \hat{x}_{ij} denote the fraction of item j assigned after. In other words, if x_{ij} is the fraction of item j assigned to agent i , we have $\bar{x}_{ij} + \hat{x}_{ij} = x_{ij}$, $\sum_j \bar{x}_{ij} v_{ij} \leq b_i$, and $\sum_j \bar{x}_{ij} v_{ij} = b_i$ if for any j , $\hat{x}_{ij} > 0$. The idea is that assigning more of item j to agent i before exceeding his budget increases both terms in the "goal" above, but assigning more after exceeding the budget only affects the second term.

⁹Traditionally, some consider only costs $c_{ij} \in \mathbb{R}_+$, but the result we cite applies for negative costs as well.

¹⁰Traditionally, it makes sense to minimize total cost. As costs are possibly negative, the use of max or min is irrelevant.

The LP relaxation is as follows:

Variables:

- \bar{x}_{ij} , for all agents i and items j , denoting the fraction of item j assigned to agent i contributing to both the budgeted-additive and additive terms in bidder i 's (virtual) welfare.
- \hat{x}_{ij} , for all agents i and items j , denoting the fraction of item j assigned to agent i contributing to just the additive term in bidder i 's (virtual) welfare.

Constraints:

- (1) $\sum_i (\bar{x}_{ij} + \hat{x}_{ij}) \leq 1$, for all j , guaranteeing that no item is allocated more than once.
- (2) $\sum_j \bar{x}_{ij} v_{ij} \leq b_i$, for all i , guaranteeing that contributions to the budgeted-additive term are not overcounted.

Maximizing:

- $\sum_{ij} m_i \bar{x}_{ij} v_{ij} + \sum_{ij} w_{ij} (\bar{x}_{ij} + \hat{x}_{ij})$, the virtual welfare. Note that as each $m_i \geq 0$ and $v_{ij} \geq 0$, the optimal solution will never have $\hat{x}_{ij} > 0$ unless $\sum_j \bar{x}_{ij} = b_i$.

It is clear that any solution to BAVWM has a corresponding fractional solution to this LP. So the goal is to solve this LP and round the fractional solution to an integral one without too much loss. The idea is that the feasible region now looks pretty similar to that of the generalized assignment problem, asking for an assignment of jobs to machines such that the capacity of machine i is at most b_i . We first prove the following rounding theorem, which is a near-direct application of Theorem 2.7.

THEOREM 3.1. *The optimal fractional solution to the above LP can be rounded in polynomial time to an integral assignment such that:*

- (1) $\sum_i (\bar{x}_{ij} + \hat{x}_{ij}) \leq 1$ for all j .
- (2) $\sum_j \bar{x}_{ij} v_{ij} \leq 2b_i$ for all i .
- (3) $\sum_{ij} m_i \bar{x}_{ij} v_{ij} + \sum_{ij} w_{ij} (\bar{x}_{ij} + \hat{x}_{ij}) \geq OPT$, where OPT is the value of the LP.

PROOF. We show how to interpret our LP as an instantiation of a fractional LP for the generalized assignment problem, and then directly apply Theorem 2.7. We use p_{ij} to denote processing times, c_{ij} to denote costs, and T_i to denote capacities in the created generalized assignment problem instance.

— **Machines:**

- (1) A dummy machine, 0.
- (2) For all bidders i , a hat machine \hat{i} (corresponding to the hat variables in our LP).
- (3) For all bidders i , a bar machine \bar{i} (corresponding to the bar variables in our LP).

— **Jobs:** A job j for all items j .

— **Processing times and costs:**

- (1) $p_{0j} = c_{0j} = 0$ for all j . $T_0 = 0$.
- (2) $\hat{p}_{ij} = 0$ for all j . $\hat{c}_{ij} = w_{ij}$ for all j . $\hat{T}_i = 0$.
- (3) $\bar{p}_{ij} = v_{ij}$. $\bar{c}_{ij} = m_i v_{ij} + w_{ij}$. $\bar{T}_i = b_i$.

The fractional LP referenced in Theorem 2.7 on this instance would then be (note that the capacity constraints for machines 0 and all \hat{i} are vacuously satisfied, and that there do not exist any i, j for which $p_{ij} > T_i$ by Observation 2):

Variables:

- x_{0j} , for all jobs j , denoting the fraction of job j assigned to machine 0.
- \bar{x}_{ij} , for all machines i and jobs j , denoting the fraction of job j assigned to machine \bar{i} .
- \hat{x}_{ij} , for all machines i and jobs j , denoting the fraction of job j assigned to machine \hat{i} .

Constraints:

- (1) $x_{0j} + \sum_i (\bar{x}_{ij} + \hat{x}_{ij}) = 1$, for all j , guaranteeing that every job is allocated exactly once.
- (2) $\sum_j \bar{x}_{ij} v_{ij} \leq b_i$, for all i , guaranteeing that the total processing time on machine \bar{i} is at most b_i .

Maximizing:

- $\sum_{ij} m_i \bar{x}_{ij} v_{ij} + \sum_{ij} w_{ij} (\bar{x}_{ij} + \hat{x}_{ij})$, the cost.

It's clear that this LP is exactly the same as our LP, just with an additional dummy bidder 0 who collects all unallocated fractions of items. By Theorem 2.7, the optimal fractional solution to this LP can be rounded in polynomial time to an integral solution whose total cost is at least as large, but where the capacity of machine \bar{i} could be as large as $2b_i$, which is exactly an integral allocation of items to bidders with the desired properties. \square

After applying Theorem 3.1, we now have an integral solution that is at least as good as the optimum, except our solution is infeasible. It's infeasible because it's "getting credit" for (virtual) welfare in the budgeted-additive term that is perhaps up to twice the budget (i.e. up to $2b_i$). An "obvious" fix to this problem might be to take this integral solution and only take credit for budgeted-additive values up to b_i , thereby making the solution feasible again. Unfortunately, because the objective is mixed sign, the resulting solution doesn't provide any approximation guarantee.¹¹ Instead, we provide a simple procedure to select a feasible suballocation of this infeasible one that loses a factor of 3.

THEOREM 3.2. *Given an integral allocation \bar{x} satisfying $\sum_i \bar{x}_{ij} + \hat{x}_{ij} \leq 1$ for all j , $\sum_j \bar{x}_{ij} v_{ij} \leq 2b_i$ for all i , and $\sum_{ij} m_i \bar{x}_{ij} v_{ij} + \sum_{ij} w_{ij} (\bar{x}_{ij} + \hat{x}_{ij}) = C$, one can find in poly-time an integral allocation \bar{y} such that:*

- (1) $\sum_i (\bar{y}_{ij} + \hat{y}_{ij}) \leq 1$ for all j .
- (2) $\sum_j \bar{y}_{ij} v_{ij} \leq b_i$ for all i .
- (3) $\sum_{ij} m_i \bar{y}_{ij} v_{ij} + \sum_{ij} w_{ij} (\bar{y}_{ij} + \hat{y}_{ij}) \geq C/3$.

PROOF. For each i , we wish to partition the set of items assigned to i via \bar{x}_{ij} (of the infeasible integral solution), S , into three disjoint sets S_i^1, S_i^2, S_i^3 such that $\sum_{j \in S_i^k} v_{ij} \leq b_i$ for all k . This is always possible: consider sorting the elements in decreasing order of v_{ij} and greedily adding them one at a time to the S_i^k with minimal weight so far. Assume for contradiction that some item j^* , when added, pushes some S_i^k from below b_i to above b_i . Then without j^* , each of S_i^1, S_i^2, S_i^3 must have had weight strictly larger than $b_i - v_{ij^*}$. As the total weight in all three (without j^*) is at most $2b_i - v_{ij^*}$, this means

¹¹Consider, for example, the following instance: there is one buyer and two items. $v_{11} = v_{12} = 3$, $b_1 = 3$, $w_{11} = w_{12} = -2$. Then the allocation that awards both items and "gets credit" for up to $2b_i$ is believed to have virtual welfare 2. However, the correctly computed virtual welfare of this allocation is actually -1 , which clearly provides no meaningful approximation. Instead we must develop a procedure that, on this instance, would allocate just one of the items.

that $2b_i - v_{ij^*} > 3(b_i - v_{ij^*}) \Rightarrow v_{ij^*} > b_i/2$. But as we processed elements in decreasing order of v_{ij} , this would imply that j^* was the third (or earlier) item processed, meaning that some set must have been empty, and j^* couldn't have possibly pushed it over the limit (as $v_{ij} \leq b_i$ for all j). Therefore, at termination we must have $\sum_{j \in S_i^k} v_{ij} \leq b_i$ for all k . Now, define $k^* = \arg \max_k \{\sum_{j \in S_i^k} m_i v_{ij} + w_{ij}\}$. Let $\bar{y}_{ij} = 1$ iff $j \in S_i^{k^*}$, and $\hat{y}_{ij} = \hat{x}_{ij}$ for all j .

It's clear that $\sum_j \bar{y}_{ij} v_{ij} \leq b_i$ for all i . As $\bar{y}_{ij} \leq \bar{x}_{ij}$ for all i, j , it's also clear that $\sum_i \bar{y}_{ij} + \hat{y}_{ij} \leq 1$ for all j . Finally, by choice of k^* it's also clear that $\sum_{ij} (m_i v_{ij} + w_{ij}) \bar{y}_{ij} \geq \sum_{ij} (m_i v_{ij} + w_{ij}) \bar{x}_{ij}/3$, and therefore $\sum_{ij} m_i \bar{y}_{ij} v_{ij} + \sum_{ij} w_{ij} (\bar{y}_{ij} + \hat{y}_{ij}) \geq C/3$, as desired. \square

Combining Theorems 3.1 and 3.2 yields a feasible, integral allocation that is a 3-approximation by rounding the fractional solution output by our LP, and it is easy to see that the entire procedure runs in polynomial time.

THEOREM 3.3. *There is a poly-time 3-approximation algorithm for Budgeted-Additive Virtual Welfare Maximization, which is a reformulation of GOOP($[n + 1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE}$). Therefore, for all $\epsilon > 0$, there is a poly-time $(\epsilon, 3)$ -approximation algorithm for BMeD($[n + 1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE}$). Specifically, if ℓ is the input length to an instance of BMeD($[n + 1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE}$), the algorithm terminates in time $\text{poly}(\ell, 1/\epsilon)$ and succeeds with probability $1 - \exp(-\text{poly}(\ell, 1/\epsilon))$.*

We conclude this section with a remark about the special case of a single (or small constant) number of items. Notice that BAVWM can be solved exactly by exhaustive search in time $\text{poly}(n^m)$. If m is a small constant, exhaustive search may be computationally feasible, resulting in an exact algorithm (instead of a 3-approximation).

Remark 3.4. Budgeted-Additive Virtual Welfare Maximization can be solved exactly in time $\text{poly}(n^m)$ by exhaustive search. Therefore, for all $\epsilon > 0$, there is an $(\epsilon, 1)$ -approximation algorithm for BMeD($[n + 1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE}$). Specifically, if ℓ is the input length to an instance of BMeD($[n + 1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE}$), the algorithm terminates in time $\text{poly}(\ell, n^m, 1/\epsilon)$ and succeeds with probability $1 - \exp(-\text{poly}(\ell, 1/\epsilon))$.

Finally, we remark that the single-item case is *especially* simpler than even the two item case. We refer the reader to [Cai et al. 2012, 2013] for complete details, but essentially the sampling procedure that results in the ϵ error of Theorem 2.5 can be replaced by an exact computation *only* in the single item case (and not even in the two item case), and ϵ can be set to exactly 0.

Remark 3.5. Budgeted-Additive Virtual Welfare Maximization with $m = 1$ can be solved exactly in time $\text{poly}(n)$ by exhaustive search: there are only n possible outcomes, corresponding to assigning the item to exactly one of the agents. Therefore, there is a $(0, 1)$ -approximation algorithm (i.e. an exact algorithm) for BMeD($[n + 1], \mathbb{R}_+^2, \text{REVENUE}$) (i.e. the single item case). Specifically, if ℓ is the input length to an instance of BMeD($[n + 1], \mathbb{R}_+^2, \text{REVENUE}$), the algorithm terminates in time $\text{poly}(\ell)$, and succeeds with probability 1.

3.2. Structural Results

In this section, we discuss the structure of the optimal mechanism, and of the computationally efficient mechanism from Section 3.1. We begin by characterizing the optimal mechanism by combining Theorem 2.6 with Observation 1.

THEOREM 3.6. *In any $\text{BMeD}([n+1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE})$ instance, the optimal mechanism can be implemented as a distribution over virtual welfare maximizers. Specifically, there exists a distribution Δ over mappings $(f_1^\delta, \dots, f_n^\delta)$. Each mapping f_i^δ maps types $(\vec{v}_i, b_i) \in \mathbb{R}_+^{m+1}$ to a multiplier $m_i^\delta(\vec{v}_i, b_i) \in \mathbb{R}_+$ and a vector $\vec{w}^\delta(\vec{v}_i, b_i) \in \mathbb{R}^m$. Define ϕ_i^δ to be the mapping that takes as input types $(\vec{v}_i, b_i) \in \mathbb{R}_+^{m+1}$ and outputs a valuation function $\phi_i^\delta(\vec{v}_i, b_i)(\cdot)$ with $\phi_i^\delta(\vec{v}_i, b_i)(S) = m_i^\delta(\vec{v}_i, b_i) \cdot \min\{b_i, \sum_{j \in S} v_{ij}\} + \sum_{j \in S} w_{ij}^\delta(\vec{v}_i, b_i)$. The allocation rule of the optimal mechanism first samples $(f_1^\delta, \dots, f_n^\delta)$ from Δ , and on profile (\vec{v}, \vec{b}) , allocates the items according to $\arg \max_{S_1 \sqcup \dots \sqcup S_n \subseteq [m]} \{\sum_i \phi_i^\delta(\vec{v}_i, b_i)(S_i)\}$. Furthermore, if $m_i^\delta(\vec{v}_i, b_i) > 0$, bidder i is charged $\min\{b_i, \sum_{j \in S_i} v_{ij}\}$. If $m_i^\delta(\vec{v}_i, b_i) = 0$, then bidder i is charged 0.*

PROOF. The proof starts with an application of Theorem 2.6 to the problem $\text{BMeD}([n+1]^m, \mathbb{R}_+^{m+1}, \text{REVENUE})$. By Observation 1, the joint optimization over allocations x and price vectors \vec{p} can be accomplished by transforming the optimization into one that depends only on the allocation. Once the allocation is found, optimization of the price vector follows as in Observation 1. \square

We remark that the virtual types involved in Theorem 3.6 have valuation functions that are the sum of a budgeted-additive function, and an additive function (the latter may have negative item values). We also note that the budgeted-additive component depends in a very structured way on the input type (\vec{v}_i, b_i) . Specifically, b_i is turned into a hard cap on the bidder's maximum valuation instead of a hard budget on her ability to pay, and the additive valuation \vec{v}_i is kept the same, forming a budgeted-additive function that is scaled by a positive multiplier m_i . The multiplier m_i and additional values \vec{w}_i may show little structure with respect to the input types (or perhaps none at all).

We also remark that the structure is especially simple in the case of a single item, because a budgeted-additive function for a single item is just a typical valuation function (where the bidder's value for the item is the minimum of her value and her budget). Specifically, the virtual type parameterized by $m_i^\delta(v_i, b_i)$ and $w_i(v_i, b_i)$ values the item at $m_i \min\{v_i, b_i\} + w_i(v_i, b_i)$. This observation leads to the following simplification:

Remark 3.7. In any $\text{BMeD}([n+1], \mathbb{R}_+^2, \text{REVENUE})$ instance (i.e. the single item case), the optimal mechanism can be implemented as a distribution over virtual value maximizers. Specifically, there exists a distribution Δ over mappings $(f_1^\delta, \dots, f_n^\delta)$. Each mapping f_i^δ maps types $(v_i, b_i) \in \mathbb{R}_+^2$ to an indicator bit $m_i^\delta(v_i, b_i) \in \{0, 1\}$ and a virtual value $\phi_i^\delta(v_i, b_i)$. The allocation rule of the optimal mechanism first samples $(f_1^\delta, \dots, f_n^\delta)$ from Δ , and on profile (\vec{v}, \vec{b}) , allocates the item to any bidder $i^* \in \arg \max_i \{\phi_i^\delta(v_i, b_i)\}$ if her virtual value is non-negative, and doesn't allocate the item otherwise. Furthermore, if $m_{i^*}^\delta(v_{i^*}, b_{i^*}) = 1$, bidder i^* is charged $\min\{b_{i^*}, v_{i^*}\}$. If $m_{i^*}^\delta(v_{i^*}, b_{i^*}) = 0$, then bidder i^* is charged 0.

We conclude with a statement regarding the format of our computationally efficient mechanisms from Section 3.1. This is an instantiation of Algorithm 2 in [Cai et al. 2013], which is used to prove Theorem 2.5.

THEOREM 3.8. *The mechanism providing the guarantee of Theorem 3.3 has the following format:*

Phase One, Find the Mechanism:

- (1) *Write a linear program that optimizes revenue over the space of truthful, feasible interim forms (Section 2.2).*

- (2) Pick an $\epsilon > 0$. Using the algorithm developed in Section 3.1, and the reduction of [Cai et al. 2013], solve this linear program approximately.
- (3) This yields an interim form corresponding to a mechanism that is an $(\epsilon, 3)$ -approximation.
- (4) The linear program also outputs auxiliary information in the form of a distribution Δ over mappings $(f_1^\delta, \dots, f_n^\delta)$ of the same format from Theorem 3.6.

Phase Two, Run the Mechanism:

- (1) Sample a mapping from Δ (provided in Phase One).
- (2) On profile (\vec{v}, b) , run the approximation algorithm of Section 3.1 for Budgeted-Additive Virtual Welfare Maximization, with input budgets b_i , input values v_{ij} , input price multipliers $m_i^\delta(\vec{v}_i, b_i)$, and input virtual values $w_{ij}^\delta(\vec{v}_i, b_i)$. Select this allocation.
- (3) If $m_i^\delta(\vec{v}_i, b_i) > 0$, charge bidder i the minimum of their budget and their value for the items they receive. Otherwise, charge them nothing.

Note that this mechanism has basically the same structure as the optimal mechanism, except that on every profile it only approximately maximizes virtual welfare (and we also first have to find the mechanism, which is completely described by the distribution Δ). In the special case of a single item, the structure can again be simplified.

Remark 3.9. In the special case of a single item, the following algorithm finds the optimal mechanism in polynomial time:

Phase One, Find the Mechanism:

- (1) Write a linear program that optimizes revenue over the space of truthful, feasible interim forms (Section 2.2).
- (2) Using the reduction of [Cai et al. 2013] and the observation in Remark 3.5 that Budgeted-Additive Virtual Welfare Maximization with $m = 1$ can be solved exactly, solve this linear program exactly. This yields an interim form corresponding to the optimal mechanism.
- (3) The linear program also outputs auxiliary information in the form of a distribution Δ over mappings $(f_1^\delta, \dots, f_n^\delta)$ of the same format from Remark 3.7.

Phase Two, Run the Mechanism:

- (1) Sample a mapping from Δ (provided in Phase One).
- (2) On profile (\vec{v}, b) , award item j to any bidder $i^* \in \arg \max_i \{\phi_i^\delta(v_i, b_i)\}$ if her virtual value is non-negative. Don't allocate item j otherwise.
- (3) If $m_i^\delta(\vec{v}_i, b_i) = 1$, charge bidder i the minimum of their budget and their value for the items they receive. Otherwise, charge them nothing.

REFERENCES

- ABRAMS, Z. 2006. Revenue maximization when bidders have budgets. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*. SODA '06. 1074–1082.
- BENOIT, J.-P. AND KRISHNA, V. 2001. Multiple-object auctions with budget constrained bidders. *Rev. Econ. Stud.*, 155–179.
- BHALGAT, A., GOLLAPUDI, S., AND MUNAGALA, K. 2013. Optimal auctions via the multiplicative weight method. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. ACM, New York, NY, USA, 73–90.
- BHATTACHARYA, S., GOEL, G., GOLLAPUDI, S., AND MUNAGALA, K. 2010. Budget constrained auctions with heterogeneous items. In *STOC*. 379–388.

- BORGS, C., CHAYES, J. T., IMMORLICA, N., MAHDIAN, M., AND SABERI, A. 2005. Multi-unit auctions with budget-constrained bidders. In *ACM Conference on Electronic Commerce*. 44–51.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2012. An algorithmic characterization of multi-dimensional mechanisms. In *the 44th Symposium on Theory of Computing Conference (STOC)*.
- CAI, Y., DASKALAKIS, C., AND WEINBERG, S. M. 2013. Understanding incentives: Mechanism design becomes algorithm design. In *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. FOCS '13. 618–627.
- CHAKRABARTY, D. AND GOEL, G. 2010. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP. *SIAM J. Comput.* 39, 6, 2189–2211.
- CHAWLA, S., MALEC, D. L., AND MALEKIAN, A. 2011. Bayesian mechanism design for budget-constrained agents. In *ACM Conference on Electronic Commerce*. 253–262.
- CHE, Y.-K. AND GALE, I. 1998. Standard auctions with financially constrained bidders. *Review of Economic Studies* 65, 1, 1–21.
- CHE, Y.-K. AND GALE, I. 2000. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic Theory* 92, 2, 198–233.
- DASKALAKIS, C. AND WEINBERG, S. M. 2015. Bayesian truthful mechanisms for job scheduling from bi-criterion approximation algorithms. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '15. 1934–1952.
- DEVANUR, N. R., HA, B. Q., AND HARTLINE, J. D. 2013. Prior-free auctions for budgeted agents. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. 287–304.
- DOBZINSKI, S., LAVI, R., AND NISAN, N. 2008. Multi-unit auctions with budget limits. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*. FOCS '08. IEEE Computer Society, Washington, DC, USA, 260–269.
- GOEL, G., MIRROKNI, V., AND PAES LEME, R. 2012. Polyhedral clinching auctions and the adwords polytope. In *Proceedings of the 44th symposium on Theory of Computing*. STOC '12. Association for Computing Machinery, 107–122.
- LAFFONT, J.-J. AND ROBERT, J. 1996. Optimal auction with financially constrained buyers. *Economics Letters* 52, 2, 181–186.
- MALAKHOV, A. AND VOHRA, R. V. 2005. Optimal Auctions for Asymmetrically Budget Constrained Bidders. Discussion Papers 1419, Northwestern University, Center for Mathematical Studies in Economics and Management Science. Dec.
- MASKIN, E. 2000. Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review* 44, 4-6, 667–681.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1, 58–73.
- PAI, M. M. AND VOHRA, R. 2014. Optimal auctions with financially constrained buyers. *Journal of Economic Theory* 150, C, 383–425.
- SHMOYS, D. B. AND TARDOS, É. 1993. Scheduling Unrelated Machines with Costs. In *the 4th Symposium on Discrete Algorithms (SODA)*.